

**Hybrid Metaheuristic Planning and Military Decision-Making:
Commonalities between theory and practice**

Topics: Lessons Learned, C2 Concepts and organizations, Social domain issues

J.L. de Jong*, T.J. Grant

POC: Tim Grant, TJ.Grant@nlda.nl

Netherlands Defence Academy

P.O. Box 90002, 4800 PA Breda, The Netherlands

tel.: +31 76 527 3261, fax: +31 76 527 3259

Abstract: The paper compares the state of the art in two different problem areas: artificial intelligence (AI) theory and the practice of military planning. In particular, hybrid metaheuristic scheduling and the operational decision-making process are compared. Although very different in nature, they share a striking number of commonalities resulting from their focus on real-world problems. This paper proposes ways in which each field could benefit from the other.

1. Introduction

The Netherlands Defence Academy (NLDA) is responsible for academic officer training for all four services (Royal Dutch Army, Air Force, Navy, and Marechaussee). In addition, it conducts research under the following six research themes:

1. Optimising the operational effort of manpower and materiel.
2. Preparing for future wars.
3. Intelligent support for operational decision-making.
4. Technology-induce transformation.
5. Collaboration between organisations.
6. Availability of military manpower and equipment.

This paper is concerned with theme (3). Key research questions in providing intelligent support for operational decision making are:

1. How can we speed up the operational decision-making process?
2. How can we ensure that the correct decisions are taken?

The research reported here focuses on speeding-up the operational decision-making process. A time-consuming part of this process is generating plans. The paper summarises the results from a preliminary study to compare military practice in generating operational plans with the state of the art in AI planning theory, with particular reference to hybrid metaheuristic scheduling. Both military practice and planning theory are undergoing transformation. Although the reasons are totally different, striking similarities are emerging. Our research was motivated by the idea that these similarities could be exploited for mutual benefit.

The paper outlines AI planning theory and hybrid metaheuristic scheduling (HMS) in particular, comparing HMS to the military operational decision-making process (ODMP). The commonalities are highlighted, and the potential implications for the military planning process are derived.

Report Documentation Page				Form Approved OMB No. 0704-0188	
Public reporting burden for the collection of information is estimated to average 1 hour per response, including the time for reviewing instructions, searching existing data sources, gathering and maintaining the data needed, and completing and reviewing the collection of information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing this burden, to Washington Headquarters Services, Directorate for Information Operations and Reports, 1215 Jefferson Davis Highway, Suite 1204, Arlington VA 22202-4302. Respondents should be aware that notwithstanding any other provision of law, no person shall be subject to a penalty for failing to comply with a collection of information if it does not display a currently valid OMB control number.					
1. REPORT DATE JUN 2006		2. REPORT TYPE		3. DATES COVERED 00-00-2006 to 00-00-2006	
4. TITLE AND SUBTITLE Hybrid Metaheuristic Planning and Military Decision-Making: Commonalities between theory and practice				5a. CONTRACT NUMBER	
				5b. GRANT NUMBER	
				5c. PROGRAM ELEMENT NUMBER	
6. AUTHOR(S)				5d. PROJECT NUMBER	
				5e. TASK NUMBER	
				5f. WORK UNIT NUMBER	
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) Netherlands Defence Academy, PO Box 90002, 4800 PA Breda, The Netherlands, ,				8. PERFORMING ORGANIZATION REPORT NUMBER	
9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES)				10. SPONSOR/MONITOR'S ACRONYM(S)	
				11. SPONSOR/MONITOR'S REPORT NUMBER(S)	
12. DISTRIBUTION/AVAILABILITY STATEMENT Approved for public release; distribution unlimited					
13. SUPPLEMENTARY NOTES The original document contains color images.					
14. ABSTRACT					
15. SUBJECT TERMS					
16. SECURITY CLASSIFICATION OF:			17. LIMITATION OF ABSTRACT	18. NUMBER OF PAGES 32	19a. NAME OF RESPONSIBLE PERSON
a. REPORT unclassified	b. ABSTRACT unclassified	c. THIS PAGE unclassified			

2. AI Planning Theory

2.1 Planning and scheduling

Planning and scheduling (P&S) problems occur everywhere in everyday life. Every situation in which actions are ordered or timed can be regarded a P&S problem (Biundo et al, 2003). Examples of P&S range through:

- *Parameter optimisation*. This involves finding the best set of parameter values. Parameters are unordered.
- *Budgeting*. This involves matching a set of resources to a set of needs, where the mapping from resources to needs is irrelevant.
- *Resource allocation*. This is similar to budgeting, but where the mapping is relevant.
- *Timetabling*. This is a specialisation of resource allocation with chronologically-ordered needs.
- *Sequencing*. This involves ordering actions into a logical sequence. The sequence may be linearly or partially (non-linearly) ordered.
- *Scheduling*. This involves ordering actions chronologically and assigning them start- and end-times.
- *Design or assembly*. This involves ordering actions in spatial order or in the order of assembly.
- *Routing or navigation*. This involves actions both spatially and chronologically.

2.2 Classical planning in AI

In AI, planning research focuses on the generation of plans, either automatically or in cooperation between automation and human planners (*mixed-initiative* planning). Plan generation is defined as the process of composing a set of operator schemata and assigning times and resources into a plan that, when executed, will change the state of the application domain from an initial state into a desired goal state. An operator schema is the generalised representation of a state-transition. Typically, the STRIPS representation (Fikes, Hart & Nilsson, 1972) is used, in which a schema has preconditions and effects. The preconditions are the elements of the domain state that must be true for the state-transition to take place, and the effects are the additions to and deletions of these state-elements that occur once the state-transition has taken place. Plan-generation algorithms exploit these preconditions and effects in selecting and sequencing the schemata.

What distinguishes AI from non-AI planning is the treatment of domain-specific knowledge. In AI planning, only the operator schemata are domain-specific. This means that the same plan-generation algorithm can – in principle - be used to generate plans for different domains just by changing the set of operator schemata, similar to using the same compiler to compile different source-code programs.

Plan generation has shown to be intractable for complex real-world problems. To make such problems tractable, restrictive assumptions must be made. The assumptions in *classical planning* are (Ghallab, Nau & Traverso, 2004):

- The application domain exhibits a finite set of states.
- All information about the application domain is available before plan generation begins, i.e. no observations need to be made to collect information about its state or the possible transitions that could occur.
- The domain is deterministic, i.e. there is no uncertainty and all information is complete and correct.
- All domain events are controllable by the planner, i.e. no external agents can change the state of the domain.
- The set of possible goal states is itself finite and restricted by comparison to the set of all possible domain states.
- All plans to be generated are sequential, i.e. linearly ordered.

- Time is implicit, i.e. the duration of states is unimportant, transitions are instantaneous, and there are no temporal constraints such as deadlines.

Real-world planning domains violate one or more – if not all - of these assumptions. For example, a domain containing consumable resources such as fuel or money will exhibit an infinite set of states. Where there are other “players” in the domain then the fourth assumption will be violated. Time is important in almost every practical problem, turning planning into scheduling.

The state of the art in AI planning research is to create algorithms for problems in which one or more of the classical planning assumptions are relaxed.

2.3 Hybrid metaheuristic scheduling

One approach to generating plans for non-classical problems is to maintain a set of algorithms. During the plan-generation process, the most suitable algorithm for the stage reached in the process is used. Several algorithms will be applied in the course of generating a plan. This approach is known as *hybrid* planning. If time or resources are involved, then it is known as hybrid scheduling.

This way of linking algorithms together is called *hybridization* (Gomes and Selman, 2001; Grant, 1986). Hybrid algorithms are a very recent playing field in metaheuristics research. Hybridization can be done in several ways. For example, different algorithms can act as islands within the bigger planning system, occasionally exchanging information. They can also be fully interleaved, resulting in a new algorithm. The algorithms can do their work in turn, resulting in a stovepipe model, as opposed to the teamwork model where the algorithms run in parallel. Talbi (2002) has developed an extensive taxonomy of hybrid metaheuristics.

As we have seen, planning problems may not turn out to be as expected from the type of problem that is to be solved. Moreover, the problem itself may change in dynamic environments. Applying a fixed hybridization may yield good results, but a really robust hybrid system should choose its constituents on the fly. Burke et al. (2003a) describe systems that pick algorithms from a portfolio. An interesting aspect of these systems is that they need an intelligent decision maker to make the actual algorithm choices (Burke et al., 2003b).

Additional problems that hybridization introduces are:

1. Which algorithm is most suitable for the current stage of the plan generation process?
2. In what order should the algorithms be applied and for how long?

In other words, hybrid planning introduces another P&S problem at a higher (*meta*) level.

The meta-level problem must be solved during the course of the plan-generation process, i.e. in real time. A common technique for controlling P&S problems in real time is to use *heuristics*, i.e. rules of thumb. Heuristics for a higher-level problem are known as *metaheuristics*.

Metaheuristics are search methods that are not problem-specific (Blum & Roli, 2003). They are more generic in nature, automatically searching solution spaces that can be described in terms of a fitness landscape. Different metaheuristics have been developed over the last two decades. Recently, they have been grouped and compared as different specimens of the same family. Evolutionary algorithms and tabu search are examples of well-known metaheuristics. All metaheuristics have in common that they perform two basic operations: they *explore* the search space and they *exploit* the promising regions. Sometimes these two processes take place in parallel, sometimes in a more sequential manner.

Different metaheuristics have different characteristics. Some put a stronger accent on exploration, others on exploitation. Depending on the problem, a certain metaheuristic may be more suitable than another. For example, when the fitness landscape is very large and contains a great number of peaks, the emphasis should lie on the exploration part of an algorithm, a task that can very well be performed by an evolutionary algorithm (Eiben & Smith, 2003). Other search spaces have a smaller amount of peaks, calling for emphasis on exploitation: where is the exact top of the peak? Tabu search (Glover & Laguna, 1997) would be an obvious choice for such a problem.

In more realistic dynamic planning problems, it is important to retain an amount of exploration and exploitation in parallel. Thus, when a solution becomes outdated, the planner can quickly catch up with a good alternative solution. Andrews and Tuson (2003) have shown that both strongly exploring and strongly exploiting algorithms have at times good characteristics in catching up on changing situations. This gives rise to the question of how both qualities of such algorithms can be used simultaneously.

Hybrid metaheuristic scheduling is the employment of rules of thumb to select and apply algorithms for generating schedules, i.e. for solving planning problems involving time and resources. According to Burke (2003a), HMS techniques:

- Are faster than single algorithms
- Provide solutions of better quality than single algorithms
- Are more robust than single algorithms
- Are more broadly applicable than single algorithms

These advantages of HMS techniques offer us what we are seeking, namely a potential way of speeding up the planning part of the ODMP.

2.4 Indications for applying HMS

There are a number of planning situations for which hybrid metaheuristic scheduling is indicated. These situations are described in the following sub-sections.

Hard versus soft constraints

Planning problems can be categorized along a wide range of variables. The more jobs that have to be planned and the more resources that can be used to execute the necessary operations, the more options the planner has to choose from. These choices are not free, as *hard constraints* have to be met. A *feasible solution* is a plan where all hard constraints have been met. *Soft constraints*, or optimality criteria, direct the search for the best solution even further. Unlike hard constraints, they cannot decide for solutions to be feasible or not. They provide a means of comparing different (feasible) solutions. Typical soft constraints are “minimize makespan”, or “maximize profit”. Constraints can even be mixed in nature, e.g. “minimize completion time (soft), but ensure it does not exceed 10 hours (hard)”. Quality of Service (QoS) levels define the minimum quality for soft constraints.

Multi-criteria optimization

In most real-life problems, multiple hard constraints have to be met, directed by multiple soft constraints. This introduces dilemmas, for example when a new candidate plan is valued higher according to one of the soft constraints, but worse on another. In such cases, instead of having just one optimal solution, some kind of multi-dimensional boundary has to be drawn consisting of *pareto-optimal* solutions. These solutions are optimal with respect to one or more of the criteria, but usually not all. T'kindt & Billaut (2002) present an elaborate overview of *multicriteria scheduling*.

NP-hard problems

When a problem is simple, e.g. when only a few tasks have to be performed on a limited set of resources, the possible solutions can simply be enumerated, after which the best one can be chosen. When the number of tasks, resources and constraints increase, the number of options to choose from increases as well. NP-hardness means that the amount of options is no longer a polynomial function of the problem size. In those situations, enumerating the options and choosing the best is no longer feasible (Garey & Johnson, 1979).

Dynamic problems

Using soft constraints as mentioned above makes it possible to value solutions. When the fitness of each solution is known, the solution space can be viewed as a fitness landscape, see Figure 1. Peaks in the landscape correspond to high-valued solutions. But problems may also be dynamic in nature: they continuously change as new orders come in and the execution of others has finished, variables or constraints change, etc. In such cases the fitness landscape is not fixed, but varies like waves on water. There will be no single solution because the optimum is always changing (Angeline, 1997).




Figure 1. An example fitness landscape.

Typically, in dynamic planning problems, another type of hard constraint comes into play: maximum planning time. This constraint is not related to the actual problem, but it is another QoS constraint put on the planner. The planner is then forced to come up with a solution within a certain amount of time.

Incomplete and uncertain knowledge

Another aspect is that the knowledge the decision maker has of the problem instance or the problem domain may be incomplete or uncertain. If so, the problem solver must fall back on his experience to assess the situation. In the dynamic planning problem described above, the future is typically unknown. The experienced problem solver will try to make predictions to assess the viability of his plans.

3. Military Operational Decision-Making

3.1 NATO ODMP

The NATO-standard operational decision-making process (ODMP) is taught to officer cadets of all military services in many nations, including the Netherlands. It is designed for hierarchical organisations, clear goals, and rational decision makers. Characteristics of the ODMP are:

- It emphasises the planning process before an operation begins.
- It assumes a hierarchical organisation.
- It is a successive, top-down decomposition process.
- Planning at each level of decomposition is largely linear; see Figure 2.
- Planning at any given level can only begin when planning is complete at the next level up. Warning orders allow some overlapping.

- Synchronisation is achieved by deconfliction at the next level up, rather than peer-to-peer. Liaison officers allow some cross-hierarchy information flow.
- There is no specific provision for multi-disciplinary collaboration.
- It assumes rational decision makers. While option selection is not exhaustive, decision makers are advised to consider three own and three enemy courses of action (COAs).
- Step 7 enables limited projection by wargaming each own COA against each enemy COA. Projection is limited because neither COA is altered by confrontation with the opposing COA.
- Decision makers may partially tailor the process. This makes compliance and interoperability dependent on shared terminology, training, culture, and infrastructure.

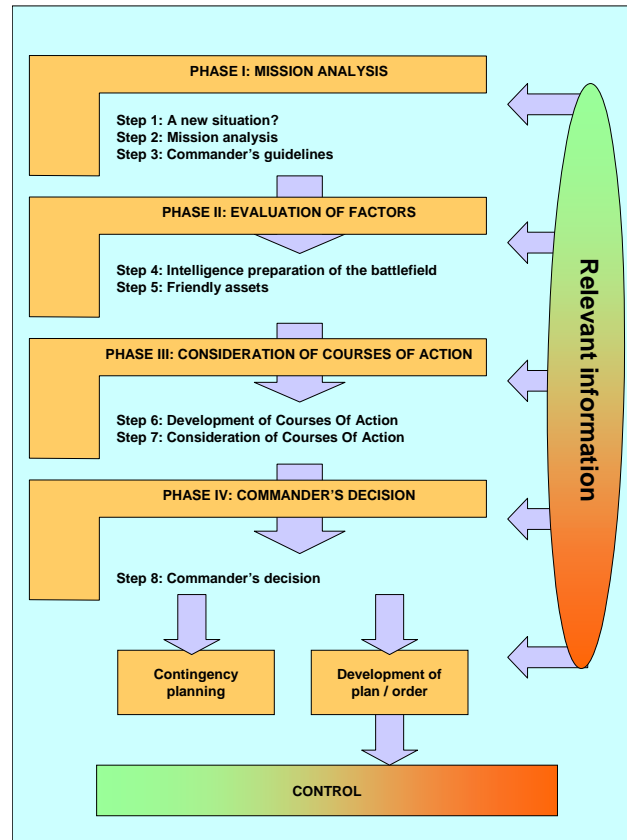


Figure 2. NATO-standard operational decision-making process.

Operational decision-making has many of the characteristics of the rational decision making model. The rational decision-maker gathers an exhaustive set of information, evaluates each option against the evaluation criteria, compares the set of options against one another, and selects the best option without regard to previous or future problems. Scientific analysis of rational decision-making centres on option evaluation and selection.

The attraction of the rational decision-making model is that it has been proven to be optimal. Its main disadvantage is that it is a time-consuming process. Moreover, optimality is dependent on the completeness and correctness of the input information. However, this is hardly ever encountered at combat situations. Moreover, the future cannot be predicted as of yet, so completeness of information can never be gained. Furthermore, the evaluation criteria should yield comparable results, otherwise no sound comparison can be made. In real life however, different criteria are of different sort, and consequently cannot be compared without the use of some translation to comparable amounts. Finally, the process assumes an unchanging world, whereas the real world actually changes during the decision making process. Enemy movements, for instance, will cause preventive movements on our side, possibly resulting in obsolete plans.

As a consequence, the theoretical guidelines are seldom followed strictly. Holewijn (2004) investigated 101 operational situations where difficult decisions had to be made, and found that the

model was properly used in only seven cases. In most cases, a straightforward decision was made based on experience. Hence, theory and practice differ greatly in most crisis response situations.

3.2 Naturalistic approach to operational decision-making

In the military domain planning problems are usually performed by human problem solvers. Research performed on planning in the military and related crisis management domains shows a number of typically recurring problem aspects. Holewijn (2004) presents a more elaborate overview of these aspects.

In real-life, decision makers often have to make up their goals as they go, as the problem is *ill defined* and *ill structured*. Apart from that, there may be *shifting goals* during execution, and there may be *competing goals*. The planner typically finds himself in an *uncertain and dynamic environment*. As a result of these dynamic goals and environments, planning cannot be a straightforward stovepipe process. Instead, *action/feedback loops* are continuously challenging decisions that have been made. *Multiple players* make the uncertainty even worse, as each has its own goals, potentially colliding with your own goal. Crisis management obviously involves a lot of *time stress* aggravated with *high stakes*.

Klein (1998) and his colleagues studied naturalistic decision-making (NDM) in a wide range of crisis management situations. They showed that expert decision-makers working in their real environments did not follow the rational decision-making model. A major incentive for them to start searching for a new decision-making model was a fire-fighter's classic remark: "I don't make decisions. I don't remember when I've ever made a decision." (Klein, 1998, p. 11).

Klein's recognition-primed decision-making (RPDM) model revolves around a number of main features. One is *the power of experience* – i.e. the experts' ability to match situations to templates developed from previous experience. A template consists of a set of cues for matching, a course of action to perform if matching succeeds, a set of expectations to check that the course of action is proceeding correctly, and a set of plausible goals. Another feature is *the power of mental simulation* – i.e. experts' ability to make a mental movie to play a course of action from start to end. In NDM, the scientific focus is on situation assessment, *not* on option evaluation and selection. The decision maker is primed to act with incomplete information, *not* to wait for complete analyses.

RPDM was developed as a descriptive approach to crisis response decision-making. Subsequently, Schmitt and Klein (1999) derived a prescriptive model of planning, known as the "Recognitional Planning Model". It features lessons learned from the previous research. For example, war-gaming is introduced to gain the experience needed to cope with difficult decision-making situations, and to rehearse the planned course of action.

4. Comparing ODM to HMS

4.1 Features for comparison

The features for comparing ODM and HMS have been extracted from Klein and Klinger's (1991, Table 1, p.17) list of features of naturalistic decision-making. Hybrid metaheuristic scheduling and the two forms of operational decision-making have been assessed against these features. The results of our evaluation are summarised in Table 1.

Table 1. Comparison of HMS and ODM against features.

Features	Hybrid Metaheuristics	NATO ODM	RPDM
<i>Goals / tasks</i>	Well-defined, shifting	Well-defined, constant, & coherent	Ill-defined, shifting
<i>Information</i>	Uncertain & incomplete	Uncertain & incomplete, but intelligence reduces	Uncertain & incomplete

		uncertainty	
Criteria	Dynamic, multiple, competing	Constant, small number, coherent	Dynamic, multiple, competing
Options	Unlimited (NP-hard)	Max. 3 COAs	Single option evaluation
Organization	No limiting cultures	Hierarchical; well defined	Organizational norms
Decision makers	Experience hard-wired into algorithms	Experience nice to have.	Experienced
Cognitive representation	Constraints	Doctrine; standard operating procedures	Templates (cues, COAs, goals, & expectations)
Learning	None	By experience	By experience and from other decision-makers' stories
Projection	No forecasting	No forecasting	Mental simulation
Environment (other agents)	Multiple algorithms; shared goals	Multiple players; competitors	Multiple players; competitors
Feedback	Continuous real-time	Next operation (not this)	Real-time
Time constraints	Severe time stress	Ample time	Time stress
Stakes	Increasing	Life & death	Life & death

4.2 Commonalities

Inspection of Table 1 shows that hybrid metaheuristic scheduling and RPDM have the following in common:

- Goals and tasks are shifting.
- Input information is uncertain and incomplete.
- Decision-making is based on experience. In HMS this experience is hard-wired into the scheduling algorithms. In RPDM it takes the form of the human decision-makers' templates.
- The environment contains multiple other agents. In HMS the other agents are the multiple algorithms. In RPDM the other agents are the competitors.
- Feedback takes place in real time.
- Decision-making must be performed under time stress.
- The stakes are very high or increasing as plan generation progresses.

4.3 Differences

Hybrid metaheuristic scheduling and RPDM also have differences. Table 1 shows that they differ in the following respects:

- Whether goals and tasks are well- or ill-defined. HMS assumes well-defined goals and tasks, but RPDM assumes that they are ill-defined.
- The number of options that are evaluated at any moment. HMS can evaluate a large number of options, limited only by the NP-hardness of the algorithms. RPDM involves single-option evaluation.
- The organisation in which decision-making takes place. There are no limiting cultures on HMS, but RPDM takes place against a backdrop of organisational norms and goals.
- The cognitive representation in HMS is (typically) constraints, whereas it is templates in RPDM.
- Current HMS techniques do not incorporate learning. By contrast, decision-makers using RPDM must be able to acquire new templates by learning from experience.
- HMS techniques lack forecasting capabilities, while RPDM depends on projecting courses of action into the future using mental simulation.

- The other agents in HMS share common goals, but in RPDM their goals conflict with the decision-maker's, resulting in competition.

5. Conclusions and Recommendations

5.1 Lessons learned: HMS

We can draw lessons for HMS techniques in self-awareness, situation awareness, situation recognition, and projection into the future.

Self-awareness

An important aspect of planning in crisis management situations is awareness of your own possibilities and competences. This awareness greatly increases your ability to deploy your assets efficiently and effectively. In analogy, the algorithms that constitute the portfolio are the assets of a hybrid planning system. Self-awareness, however, is a topic that can hardly be found in AI planning literature. In the hybrid case, self-awareness would mean that the system assesses the qualities of the constituting algorithms. When can they best be deployed, when not? Which algorithms work well together, which ones do not? Which algorithms will guarantee me a certain quality of service level? Future research should investigate how the qualities of constituting algorithms can best be assessed or quantified, enabling the managing system to better deploy them.

Situation Awareness

Even more difficult than getting Self Awareness is obtaining good Situation Awareness. The vast amounts of literature available on this topic emphasize its importance in the planning field (Endsley, 2000). Knowing the, possibly changing, situation in which you find yourself enables you to better deploy yourself appropriately. Again, in AI planning literature, this virtue is underrated. Putting algorithms together to solve the problem, and reconfiguring them to suit the specific (changing) situation is not a topic of interest, as of yet. The benefits are clear: choosing algorithms in accordance to the problem could greatly increase the optimality for a specific problem.

Recognition

As section 3.4 shows, a lot of decisions are made not so much based on choosing from a large set of possibilities. Instead, expert decision makers make shortcuts from situation to action, based on fast recognition of the situation. In nearly all of AI Planning literature, however, the planner starts from scratch and completely has to re-invent the wheel every time it is run. Building a library of templates that have only to be refined in order to suit a specific problem could increase the planning speed enormously. Finding out how experience and templates are gathered is a very interesting and important topic of research.

Projection

Being able to 'play the movie in your head' helps the decision maker to evaluate his plans by mental simulation. It involves 'looking in the future' and assessing 'how the plan will work out'. This situation forecasting is a much-underrated topic in the AI planning literature. Actually, a forecast of the situation in the future should be a standard evaluation criterion in all dynamic planners, as change is a constant in these environments.

5.2 Lessons learned: ODMP

We can draw lessons for ODMP in multi-disciplinary co-operation and collaborative planning.

Multi-disciplinary co-operation

The strong point of HMS systems is the fact that they combine the strength of multiple different algorithms. Different algorithms have different strong points; each has a typical 'view' on the problem. In the military situation, bringing decision-makers from multiple disciplines around the same table is not common. A typical example would be to change the 'people portfolio' according

to the situation, and have, for example, an expert of a certain culture when performing military operations abroad. Keus (2002) initiated such research with his framework for cooperative decision making in teams.

Collaborative planning

Another form of multidisciplinary cooperation would be to let different levels within an army organization collaborate with one another. In good hybrid systems, ‘higher’ algorithms can benefit from ideas thrown up by ‘lower’ ones. By analogy, information on an emerging plan in military decision-making should be communicated both top-down and bottom-up. Computerized planning aids would enable this way of planning, where all of the different planners are planning concurrently. Such collaborative planning should dramatically reduce planning time. Furthermore, all groups involved have a better picture of the plan as a whole, and therefore a better situation awareness. Within a brigade, for example, one discipline can quickly catch up with a changing situation because it has good knowledge of what the adjacent disciplines are doing. It will also prevent higher levels from creating plans that are hard to implement at a lower level.

5.3 Recommendations

We recommend that research should be done into incorporating:

- Self-awareness, situation awareness, situation recognition, and projection in hybrid metaheuristic scheduling.
- Multi-disciplinary co-operation and collaborative planning in the operational decision-making process.

6. Acknowledgements

This study was carried out at and funded by the Royal Netherlands Defence Academy (<http://www.nlda.nl>). The hybrid metaheuristics scheduling part of this work was also supported by the Dutch Ministry of Economic Affairs, grant no: BSIK03024, within the ICIS project. The ICIS project is hosted by the DECIS Lab (<http://www.decis.nl>), an open research partnership of Thales Nederland, the Delft University of Technology, the University of Amsterdam and the Netherlands Foundation of Applied Scientific Research (TNO).

7. References

- Angeline, P. (1997) *Tracking Extrema in Dynamic Environments*. Proceedings of the Sixth Annual Conference on Evolutionary Programming, P.J. Angeline, R.G. Reynolds, J.R. McDonnell, and R. Eberhart (eds.) (Springer-Verlag), pp. 335-345.
- Biundo, S. Aylett, R., Beetz, M. Borrajo, D., Cesta, A., Grant, T.J., McCluskey, L., Milani, A., & Verfaillie, G. (eds.) (2003) *Technological Roadmap on Planning and Scheduling*. European Union Network of Excellence in AI Planning (PLANET), downloadable from <http://www.planet-noe.org>.
- Blum, C., and A. Roli. (2003) *Metaheuristics in combinatorial optimization: Overview and conceptual comparison*. ACM Comput. Surv. 35, 3 (Sep. 2003), 268-308.
- Burke, E.K., G. Kendall, J. Newall, E. Hart, P. Ross, and S. Schulemburg. (2003a) *Hyperheuristics: an Emerging Direction in Modern Search Technology*. In: F.W. Glover, G.A. Kochenberger (eds.), *Handbook of Metaheuristics*, Kluwer Academic Publishers, 2003.
- Burke, E.K., Kendall, G., and E. Soubeiga. (2003b) *A Tabu-Search Hyperheuristic for Timetabling and Rostering*. Journal of Heuristics, Volume 9, Issue 6, Dec 2003, pp. 451 – 470.
- Dutch Army. (2000) *Leidraad Commandovoering*. Army Field Manual I, Doctrine Commission of the Royal Dutch Army, The Netherlands. (In English: *Command & Control*)
- Dyer, D., S. Cross, C.A. Knoblock, S. Minton, and A. Tate. (2005) *Guest Editors' Introduction: Planning with Template*. IEEE Intelligent Systems, vol.20, no.2, pp. 13- 15.

- Eiben, A.E., and J.B. Smith. (2003) *Introduction to evolutionary computing*. Springer-Verlag.
- Endsley, M.R. (2000) *Theoretical underpinnings of Situation Awareness*. In M.R. Endsley and D.J. Garland. (eds), *Situation Awareness Analysis and Measurement*. LEA, Mahwah, NJ, USA.
- Fikes R.E., Hart, P.E., and Nilsson, N.J. (1972) *Learning and Executing Generalised Robot Plans*. Artificial Intelligence Journal, 3, pp. 251-288.
- Garey, M.R., and D.S. Johnson. (1979) *Computers and intractability: A guide to the theory of NP-Completeness*. W.H. Freeman.
- Glover, F., and M. Laguna. (1997) *Tabu Search*. Kluwer Academic Publishers.
- Ghallab, M., D. Nau, and P. Traverso. (2004) *Automated Planning: Theory and practice*. Morgan Kaufman publishers, San Francisco, California, USA. ISBN 1-55860-856-7.
- Gomes, C.P., and Bart Selman. (2001) *Algorithm Portfolios*. Artificial Intelligence 126, pp. 43-62.
- Grant, T.J. (1986) *Lessons for OR from AI: A Scheduling Case Study*. Journal of Operational Research Society, Vol 37:1, pp. 41-57. Opl. Res. Soc. Ltd., UK.
- Holewijn, B.J. (2004) *Snel Geschoten is Vaak Raak: Een onderzoek naar crisisbesluitvorming en het gebruik van RPD door militaire commandanten onder operationele omstandigheden*. MSc Thesis, Vrije Universiteit Amsterdam, The Netherlands. (In English: *A Fast Shot is Often On Target: Research into crisis decision-making and the use of RPD by military commanders in operational circumstances*)
- Keus, H.E. (2002) *A Framework for Analysis of Decision Processes in Teams*. Proceedings, CCRP Symposium, June 2002, Monterey, CA, USA.
- Klein, G. (1998) *Sources of power: How people make decisions*. MIT Press, Cambridge MA, second edition.
- Klein, G. and D. Klinger. (1991) *Naturalistic Decision Making*. Humans Systems IAC Gateway, XI, 3, pp. 16-19.
- Klein, M. (1991): “*Supporting conflict resolution in cooperative design systems*”. IEEE Transactions on Systems, Man and Cybernetics, vol.21, no.6, pp.1379-1390.
- Schmitt, J.F., and G.A. Klein. (1999) *A Recognition Planning Model*. Proceedings, CCRTS Symposium, Newport, RI, USA.
- Talbi, E.-G. (2002) *A Taxonomy of Hybrid Metaheuristics*. Journal of Heuristics, 8: 541-654. Kluwer Academic Publishers.
- T'kindt, V., and J-C. Billaut. (2002) *Multicriteria Scheduling - Theory, Models and Algorithms*. Springer Verlag.

Netherlands Defence Academy

Hybrid Metaheuristic Planning & Military Decision-Making

Jeroen L. de Jong

`Jeroen.deJong@decis.nl`

Prof. dr. Tim J. Grant

`TJ.Grant@nlda.nl`

Outline

Introduction

AI planning theory:

- Classical planning
- Hybrid metaheuristic scheduling

Military operational decision-making process:

- NATO process
- Naturalistic process

Comparing DMP to HMS

Conclusions & recommendations

NLDA

The NLDA logo consists of the text "NLDA" in a white, sans-serif font, positioned to the left of a stylized graphic. The graphic is composed of several parallel, diagonal lines in shades of blue and orange, creating a sense of motion or a stylized arrow pointing upwards and to the right.

AI Planning theory

Seen everywhere in everyday life:

- Resource allocation
- Timetabling
- Sensor management
- Routing or Navigation

Classic Planning (2)

Features of Classical planning problems:

- Finite set of states in the domain
- Complete & Certain knowledge of domain
- The domain is completely controllable
- Well defined goals & actions

The NLDA logo consists of the text "NLDA" in a white, sans-serif font, positioned to the left of a stylized graphic. The graphic is composed of several parallel, diagonal lines in shades of blue and orange, creating a sense of motion or a stylized arrow pointing upwards and to the right.

NLDA

Hybrid Metaheuristic Scheduling

Features of realistic planning problems:


- Infinite set of states
- Incomplete and uncertain knowledge
- The domain is uncontrollable.
- Response to own actions can be unpredictable.
- Ill-defined and competing goals

The NLDA logo consists of the text "NLDA" in a white, sans-serif font, positioned to the left of a stylized graphic. The graphic is composed of several parallel, diagonal lines in shades of blue and orange, creating a sense of motion or a stylized arrow pointing upwards and to the right.

NLDA

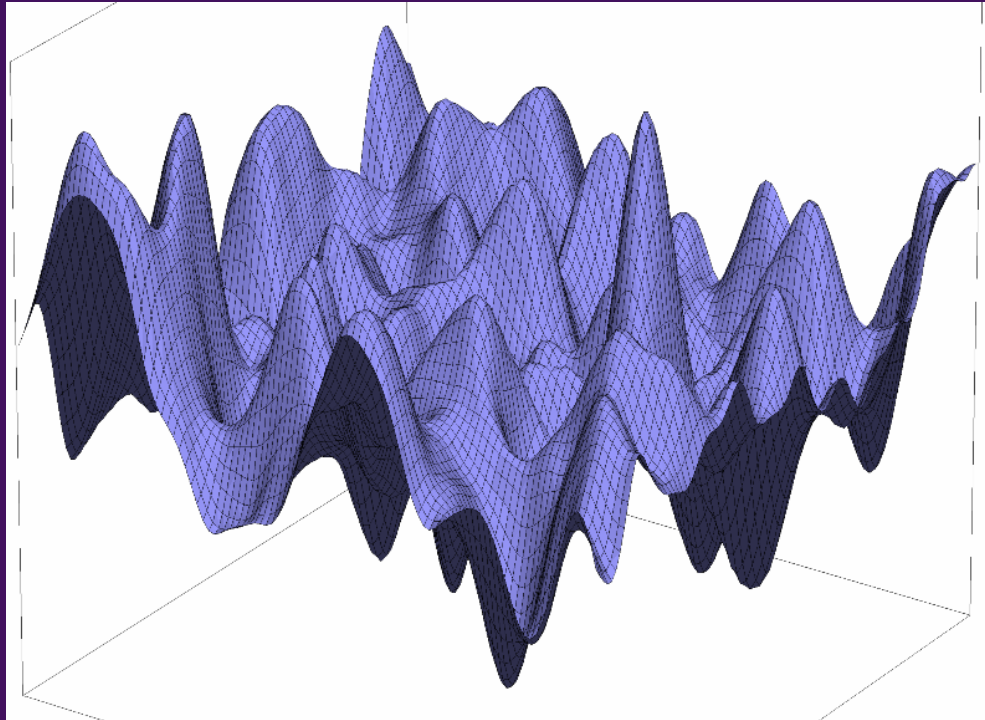
Hybrid Metaheuristic Scheduling (2)

What is hybridization?



Hybrid Metaheuristic Scheduling (3)


How does hybridization work for *Metaheuristics*?



Consider the ‘Solution Space’ as a Fitness Landscape

Hybrid Metaheuristic Scheduling (3)

How does hybridization work for *Metaheuristics*?




A single, simple heuristic finds the top

NLDA



Hybrid Metaheuristic Scheduling (3)

How does hybridization work for *Metaheuristics*?




A Metaheuristic finds promising regions



Hybrid Metaheuristic Scheduling (3)

How does hybridization work for *Metaheuristics*?



A *hybrid* metaheuristic combines best of both

NLDA

Hybrid Metaheuristic Scheduling (4)

Hybrid Algorithms (Burke 2003):

- Faster than single algorithms
- More robust
- More broadly applicable
- Better solutions

Questions for Hybrid algorithms:

- What algorithms are most suitable for the current stage of the plan generation process?
- How can they be 'glued' together?



Military decision making process

NATO Standard DMP:

1. Analyze the assignment
2. List up all possible solutions
3. Make a correct and complete list of evaluation criteria
4. Evaluate solutions according to the criteria
5. Choose the best one

Military decision making process (2)

Problems with NATO Standard DMP:

- Enormous amount of possible actions
- Time consuming
- Assumes correct & complete knowledge
- Evaluation criteria often incomparable

Therefore: hardly ever followed (Holewijn, 2004)



Military decision making process (3)

Recognition Primed Decision Making

“I don’t make decisions. I don’t remember when I’ve ever made a decision.”

-Firefighter in interview (Klein, 1998).

The NLDA logo consists of the text "NLDA" in a white, sans-serif font, positioned to the left of a stylized graphic. The graphic is composed of several parallel, diagonal lines in shades of blue and orange, creating a sense of motion or a stylized arrow pointing upwards and to the right.

NLDA

Military decision making process (4)

Recognition Primed Decision Making

- Power of intuition
- Power of mental simulation
- Focus on *situation assessment*, not on *decision events*.
- Primed to *act*, not to *wait for complete analyses*

NLDA

Comparing HMS to DMP

Features	Hybrid Metaheuristics	NATO DMP	RPDM
Goals	Well-defined, but shifting	Well-defined, constant, and coherent	Ill-defined, shifting
Information	Uncertain and incomplete	Uncertain & incomplete, but intelligence process to make it certain & complete	Uncertain, ambiguous, incomplete
Conditions	Dynamic, multi competing criteria	Constant, small number of coherent criteria	Dynamic, multi competing criteria
Outside world	Uncontrollable	Completely controllable or predictable.	Uncontrollable
Feedback	Action-feedback loops (real-time reactions to changed conditions)	To next operation, not this one	Action-feedback loops (real-time reactions to changed conditions)
Time constraints	Severe time stress	Ample time	Time Stress
Experience of Decision Makers	Experience is 'hard wired' in algorithms	Experience nice to have.	Experienced
DM: Simulation	No forecasting of future	No forecasting of future	Mental simulation to forecast plan results

NLDA

Comparing HMS to DMP (2)

Lessons from DMP to Hybrids:

- More **self awareness** – assessing the qualities of constituting algorithms
- More **Situation Awareness**– assessing the situation better, applying yourself accordingly
- **Stop starting from scratch**, do more ‘Recognition Primed’ planning – use a template when a previous problem is recognized.
- **Predict the future** – to better evaluate plans

The NLDA logo consists of the text "NLDA" in a sans-serif font, positioned above a stylized graphic of several parallel, slightly curved lines in shades of blue and orange, suggesting a dynamic or technological theme.

NLDA

Comparing HMS to DMP (3)

Lessons from HMS to Military Decision Making:

- **Multidisciplinary cooperation** – perform better in unfamiliar situations
- **Distributed planning** – perform simultaneous planning independently, afterwards evaluation & merging
- **Intertwined planning** – Different hierarchy levels and different army disciplines work on the same planning at the same time. This speeds up the process and enhances situation awareness.

The NLDA logo is located in the bottom right corner of the slide. It consists of the text "NLDA" in a white, sans-serif font, positioned above a stylized graphic of several parallel, diagonal lines in shades of blue and orange, suggesting a dynamic or technological theme.

NLDA

Summary

- Indications for applying Hybrid Metaheuristics are similar to indications that led to Recognition Primed Decision Making.
- On the solution side: Hybrid Metaheuristics are lacking Situation Awareness, Recognition skills, Learning.
- On the solution side: Military Decision Making lacks Intertwined, simultaneous planning.
- Research projects at NLDA and DECIS are aimed at answering these questions.



References

- Burke, E.K., G. Kendall, J. Newall, E. Hart, P. Ross, and S. Schulemburg. (2003a): “*Hyper-heuristics: an Emerging Direction in Modern Search Technology*”. In: F.W. Glover, G.A. Kochenberger (eds.), *Handbook of Metaheuristics*, Kluwer Academic Publishers, 2003
- Holewijn, B.J. (2004): “*Snel geschoten is vaak raak; een onderzoek naar crisisbesluitvorming en het gebruik van RPD door militaire commandanten onder operationele omstandigheden*”. MSc Thesis, Vrije Universiteit Amsterdam, The Netherlands.
- Klein, G. (1998): “*Sources of power; How people make decisions*”. MIT Press, Cambridge MA, second edition.

Questions

